Helios™ Context Continuity Architecture (HCCA™)

Whitepaper v1.0 — Codex-Centric Memory and Rules Continuity

Author / Originator: Anthony

Brand: Helios™

Version: v1.0 (Formatting Revision A)

Date: 2025

Author’s Note

This architecture emerged from extended, hands-on experimentation with large language models across long conversations and separate sessions, where continuity of memory, rules, and user preferences proved essential but unreliable when left to platform-managed mechanisms. Modern tools, including AI systems, were used to assist in drafting and refinement; however, the architectural decisions, scope boundaries, and claims presented here reflect direct human design choices and empirical observation. Responsibility for all claims, limitations, and interpretations rests solely with the author.

Abstract

Large language models (LLMs) exhibit fundamental continuity limitations: information loss across long conversations due to finite context windows and loss of state between separate conversation threads. Vendor-managed memory features, where available, are typically opaque, non-portable, and platform-dependent. This paper introduces Helios™ Context Continuity Architecture (HCCA™) v1.0, a codex-centric approach that demonstrates how memory and rules, including user preferences, can persist across context windows and across separate conversation threads through the deliberate creation and reuse of Codexes—structured artifacts that externalize saved data and rulesets. HCCA™ v1.0 establishes a minimal, explicit, and proven architecture for continuity that is independent of vendor-managed memory systems.

1. Introduction

Large language models are fundamentally stateless systems. Any apparent memory is typically confined to the active context window and is lost when that window is exceeded or when a new conversation thread is created. This behavior undermines long-running projects, stable system identity, and consistent application of user preferences and rules.

2. Scope and Claims (v1.0)

HCCA™ v1.0 demonstrates three proven capabilities: memory and rules continuity across context windows, memory and rules continuity across separate conversation threads, and the creation and reuse of Codexes as structured artifacts containing saved data and rulesets. No other mechanisms are claimed.

3. Definitions

Memory refers to persisted information required for continuity. Rules refer to constraints and preferences governing behavior. A Codex is a structured, persistent artifact that stores memory and rulesets in explicit form.

4. Codex-Centric Architecture

In HCCA™ v1.0, the Codex is the sole intentional continuity mechanism. It is created, persisted independently of conversations, and reintroduced to restore continuity when needed.

5. Limitations

Reintroduction of Codexes may be manual, and behavioral variance across models may occur. No claims are made regarding automation, enforcement, or determinism.

6. Conclusion

HCCA™ v1.0 demonstrates that memory and rules continuity—including user preferences—can persist across context windows and conversation threads through explicit Codex creation and reuse. This establishes a minimal, defensible foundation for AI continuity.